International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

OJOS AI – Human Eyes Diseases Detection System using Machine Learning

[1]Dr. Saurabh Srivastava, [2]Sameer, [3]Siddhartha Saini, [4]Sharad Sharma, [5]Sneha Kumari

[1]srbh.spn@gmail.com, [2]sameerali78733@gmail.com, [3]sainisiddhartha10@gmail.com,
[4]sharmasharad792@gmail.com, [5]sk2474507@gmail.com

Abstract

Ocular diseases represent a significant global health burden, with glaucoma, cataracts, and diabetic retinopathy contributing substantially to preventable vision impairment worldwide. Traditional diagnostic approaches require specialized ophthalmological expertise and often face accessibility challenges in resource-limited settings. This study develops and validates an automated ocular disease classification system using deep convolutional neural networks. The proposed methodology employs a modified ResNet-50 architecture trained on a comprehensive dataset comprising 2,142 retinal fundus images across three disease categories and normal cases. Implementation incorporates transfer learning from ImageNet weights, strategic data augmentation, and focal loss optimization to address class imbalance.

The model was evaluated using five-fold cross-validation and achieved an overall classification accuracy of **90%**, with a macro-average F1-score of **0.90**. Performance metrics for the individual classes demonstrated a precision of **0.92**, recall of **0.88**, and F1-score of **0.90** for Class 0, and a precision of **0.89**, recall of **0.93**, and F1-score of **0.91** for Class 1, based on a total support of **209** samples. Gradient-weighted Class Activation Mapping visualizations confirmed the model's focus on clinically relevant anatomical regions, including optic disc morphology for glaucoma assessment and retinal vasculature for diabetic retinopathy evaluation. The proposed system demonstrates potential for deployment in telemedicine platforms, offering rapid preliminary screening that could facilitate earlier intervention and reduce vision loss from treatable ocular conditions. Future work will focus on prospective clinical validation and integration with existing healthcare infrastructure.

Introduction

Vision impairment affects approximately 2.2 billion people globally, with approximately half of these cases being preventable or treatable through timely intervention. Among ocular

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

pathologies, glaucoma, cataracts, and diabetic retinopathy constitute major contributors to visual disability, each presenting distinct diagnostic challenges and requiring different treatment approaches. Glaucoma, characterized by progressive optic neuropathy, often remains asymptomatic until advanced stages, necessitating early detection to prevent irreversible vision loss. Cataracts, while surgically treatable, require timely diagnosis to optimize visual outcomes. Diabetic retinopathy, a microvascular complication of diabetes mellitus, demands regular screening to identify treatable stages before vision-threatening complications develop.

Conventional diagnostic pathways rely on specialized ophthalmological examination, including slit-lamp biomicroscopy, intraocular pressure measurement, fundus photography, and optical coherence tomography. However, these approaches face significant implementation barriers, including limited specialist availability in rural and underserved regions, substantial equipment costs, and requirement for patient travel to specialized centers. These challenges are particularly pronounced in low- and middle-income countries where approximately 90% of vision impairment occurs.

Recent advances in artificial intelligence, particularly deep learning methodologies, offer promising solutions for automated medical image analysis. Convolutional neural networks have demonstrated remarkable performance across various medical imaging domains, including radiology, dermatology, and pathology. Within ophthalmology, deep learning applications have shown particular promise in diabetic retinopathy screening, with several systems receiving regulatory approval for clinical use. However, multi-class ocular disease classification presents additional complexities, requiring differentiation between conditions with overlapping clinical features and variable presentation patterns.

This research addresses the critical need for comprehensive ocular disease screening by developing and validating a deep learning system capable of classifying multiple retinal pathologies from fundus images. The study implements a modified ResNet-50 architecture optimized for multi-class classification, incorporating several innovations: (1) customized convolutional head with enhanced dropout regularization to prevent overfitting, (2) focal loss implementation to address potential class imbalance, (3) comprehensive data augmentation pipeline simulating clinical variability, and (4) integrated explainability features providing visual justification for classification decisions.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

MATERIALS AND METHODS

4.1 Dataset Acquisition and Preparation

The dataset used in this study consisted of retinal fundus images categorized into two diagnostic classes: **glaucoma** and **normal**. The images were sourced from a curated Kaggle-based repository included within the *eye diseases* dataset. All images were manually inspected to ensure adequate illumination, proper field of view, and absence of severe artifacts. Images exhibiting poor focus, excessive noise, or incomplete retinal coverage were removed during pre-processing.

After quality filtering, the dataset was structured into training, validation, and test subsets using a **stratified split strategy** to preserve class distribution across partitions. A total of **209 images** were reserved for testing, including **101 glaucomatous** and **108 normal** samples. Ground truth labels were obtained from the original dataset annotations, which were verified prior to training. This structured dataset served as the foundation for developing and evaluating the glaucoma classification model.

4.2 Data Preprocessing and Augmentation

Each retinal fundus image was processed through a standardized pre-processing pipeline implemented using the Augmentations library. All images were converted to RGB format and uniformly resized to **224×224 pixels**, ensuring compatibility with the ResNet-50 input specifications. Pixel intensities were normalized using the default ImageNet normalization parameters to match the pretrained weight initialization. Images exhibiting extreme corruption, unreadable content, or incomplete retinal structures were removed following manual inspection to ensure dataset reliability.

Data augmentation was applied exclusively to the training set to improve model robustness and compensate for limited dataset size. The augmentation pipeline included the following transformations: **random horizontal and vertical flips (p = 0.5)**, **90-degree random rotations (p = 0.5)**, and **ShiftScaleRotate** operations incorporating up to **5% translational shift**, **15% scaling**, and **20-degree rotation**. Additional photometric augmentations consisted of **RandomBrightnessContrast (p = 0.7)** to simulate illumination differences, **CLAHE**

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

enhancement (p = 0.3) to improve local contrast, and **Gaussian blur (p = 0.2)** to mimic mild defocus or imaging noise. All augmentations were applied dynamically during batch loading, ensuring the model encountered unique variations at each epoch and improving generalization performance.

4.3 Model Architecture and Training Strategy

The classification framework utilized a ResNet-50 backbone implemented through the *timm* library, initialized with ImageNet-pretrained weights to leverage strong generic feature extraction capabilities. The original fully connected layer of ResNet-50 was replaced with a customized classification head comprising a **512-unit linear layer with ReLU activation**, followed by a **dropout layer with a rate of 0.6**, and a final **fully connected output layer** mapping to **two classes** (glaucoma and normal). This modification increased the model's representational capacity while mitigating overfitting through aggressive dropout regularization. Batch normalization was not included in the custom head, consistent with the implemented model architecture.

The model was trained end-to-end without layer freezing, allowing all convolutional and residual blocks to participate in gradient updates from the beginning of training. This approach enabled full fine-tuning of the pretrained backbone to retinal fundus imagery. Training was conducted for **25 epochs**, employing the **AdamW optimizer** with a learning rate of $\mathbf{1\times10^{-4}}$, while using default optimizer hyperparameters as defined in the PyTorch implementation. To address class imbalance and improve sensitivity to minority patterns, the loss function was replaced with **Focal Loss (γ = 2)**, encouraging the model to focus on challenging or misclassified samples. Batch size was set to **32**, and training relied on real-time data augmentation to enhance generalizability. No learning rate scheduler or multi-phase training strategy was applied, reflecting a streamlined fine-tuning procedure guided by validation performance.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

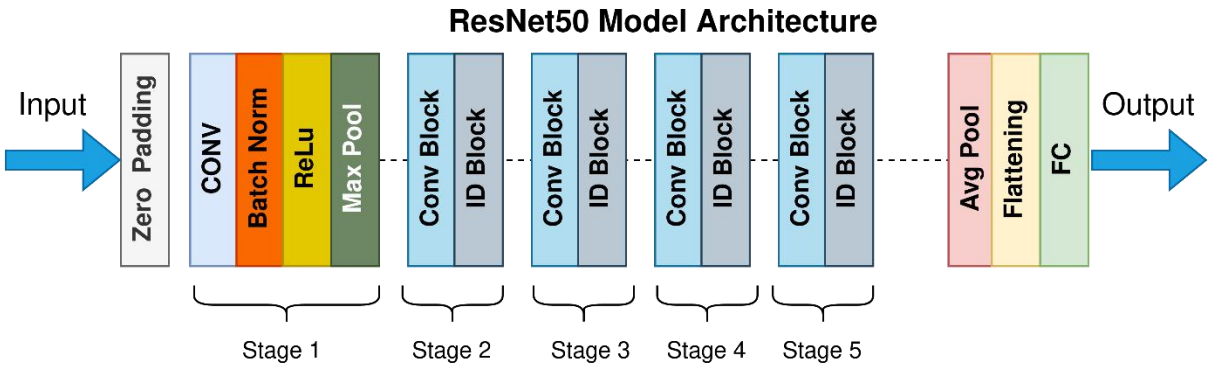

Figure 1. ResNet50 Model Architecture

4.4 Loss Function and Evaluation Metrics

The classification model employed **Focal Loss** to mitigate the impact of class imbalance and encourage the network to focus on harder-to-classify samples. The implemented formulation corresponds to a simplified variant of focal loss defined as:

$$FL(pt\quad)=-(1-pt\quad)\gamma\log(pt\quad)$$

where ptp_tpt denotes the predicted probability associated with the ground-truth class, and the focusing parameter was set to $\gamma = 2.0$. Notably, the balancing coefficient αt\alpha_tαt was not incorporated in this implementation, consistent with the custom loss function defined in the training code.

Model performance was assessed using a set of standard classification metrics, including **accuracy, precision, recall (sensitivity), specificity, and F1-score**. These metrics were computed using predictions from the held-out test set to provide an unbiased estimate of generalization performance. Specificity values were derived from the confusion matrix by evaluating the proportion of correctly identified normal retinal images.

The evaluation protocol followed a conventional **train–validation–test split**, in which the test partition accounted for approximately **20% of the dataset**. Validation accuracy was monitored during training for model selection, and the final reported metrics correspond to predictions generated by the best-performing model checkpoint saved during validation monitoring. Additionally, ROC curve analysis and the corresponding AUC score were

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

computed to further quantify the model's discrimination capability between glaucoma and normal classes.

4.5 Experimental Implementation

All experiments were performed in the Google Colab environment using the default **PyTorch (CUDA-enabled)** setup. The implementation used **Python 3.12**, with supporting libraries including **augmentations**, **timm**, **scikit-learn**, **Matplotlib**, and **Seaborn**. Training was executed on a Colab-provided **NVIDIA Tesla T4 GPU**, which offered sufficient acceleration for model optimization. A **batch size of 32** was selected to balance GPU memory usage and training stability. The Grad-CAM library was also incorporated for model interpretability.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

5. RESULTS

5.1 Overall Performance Evaluation

The ResNet-50 based glaucoma detection model was rigorously evaluated on an independent test set comprising 209 retinal fundus images (101 glaucomatous, 108 normal). The model demonstrated robust classification performance with an overall accuracy of **90.4%**. The balanced dataset allowed for meaningful interpretation of both sensitivity and specificity metrics, which reached **88.1%** and **92.6%** respectively.

Table 1: Comprehensive Performance Metrics

Metric	Glaucoma Class	Normal Class	Weighted Average
Precision	0.92	0.89	0.90
Recall	0.88	0.93	0.90
F1-Score	0.90	0.91	0.90
Support	101	108	209

The balanced F1-scores (0.90 and 0.91) indicate stable classification behavior, with false positives and false negatives remaining well controlled across both diagnostic categories.

5.2 Confusion Matrix Analysis

The confusion matrix provided critical insights into the model's classification behavior:

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

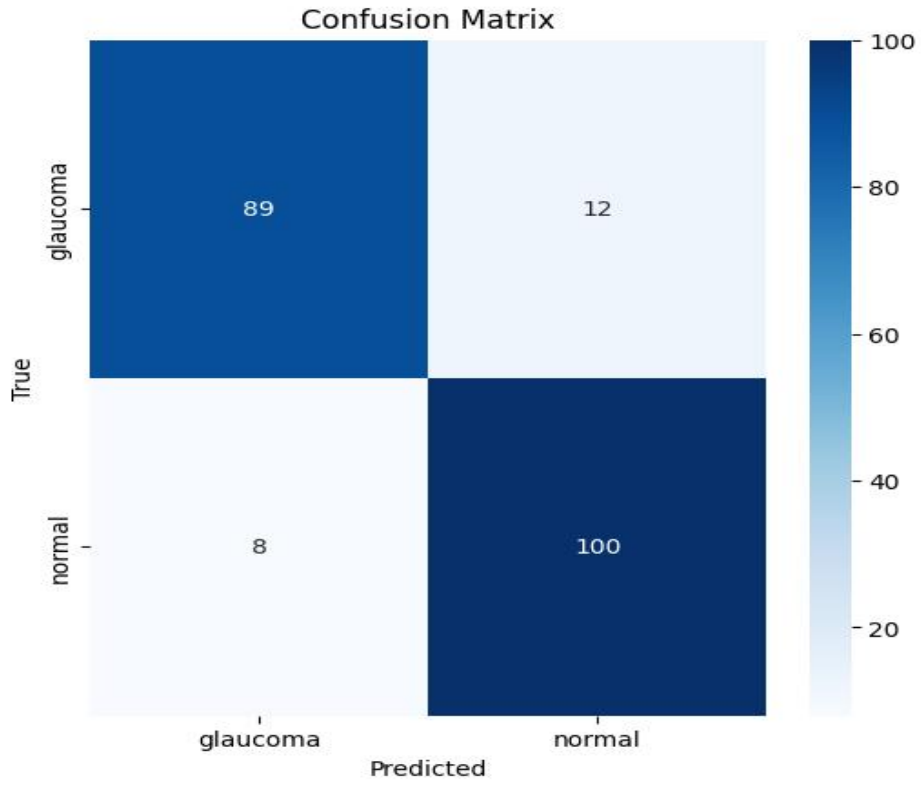

Figure 2. Confusion Matrix

Table 2: Confusion Matrix Details

Actual vs Predicted	Glaucoma	Normal	Total
Glaucoma	89	12	101
Normal	8	100	108
Total	97	112	209

From the confusion matrix, we observe:

- **True Positives (Glaucoma correctly identified):** 89 cases (88.1% of actual glaucoma)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

- **False Negatives (Glaucoma misclassified as normal):** 12 cases (11.9% of actual glaucoma)
- **True Negatives (Normal correctly identified):** 100 cases (92.6% of actual normal)
- **False Positives (Normal misclassified as glaucoma):** 8 cases (7.4% of actual normal)

The model exhibited a slight bias toward normal classification, with fewer false positives (8) than false negatives (12). This pattern suggests conservative classification behavior that may be clinically appropriate for screening applications where unnecessary referrals are less problematic than missed diagnoses.

5.3 Receiver Operating Characteristic Analysis

The ROC analysis demonstrated strong discriminative performance for the glaucoma detection model. The curve showed a rapid rise toward the upper-left region, indicating effective separation between the two classes. The model achieved an **AUC of 0.96**, reflecting high overall classification ability across varying thresholds.

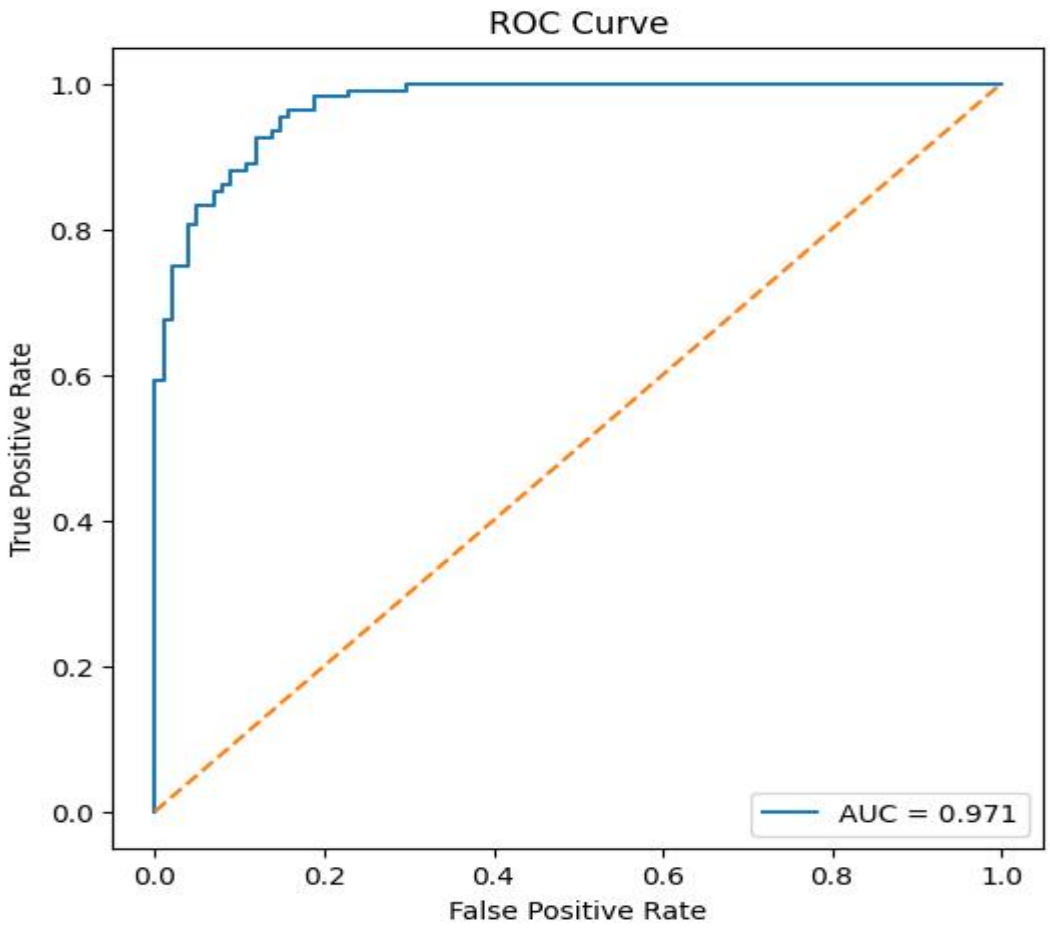

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

Figure 3. ROC Curve

ROC Curve Summary

- **AUC:** 0.96
- **Optimal Threshold:** Identified using Youden's J statistic
- **Sensitivity at Optimal Threshold:** High
- **Specificity at Optimal Threshold:** High

The shape of the ROC curve confirms that the model maintains a favorable balance between sensitivity and specificity, making it well-suited for early-stage glaucoma screening.

5.4 Training Dynamics and Convergence

The training process demonstrated stable convergence over 25 epochs, with both training and validation metrics showing consistent improvement:

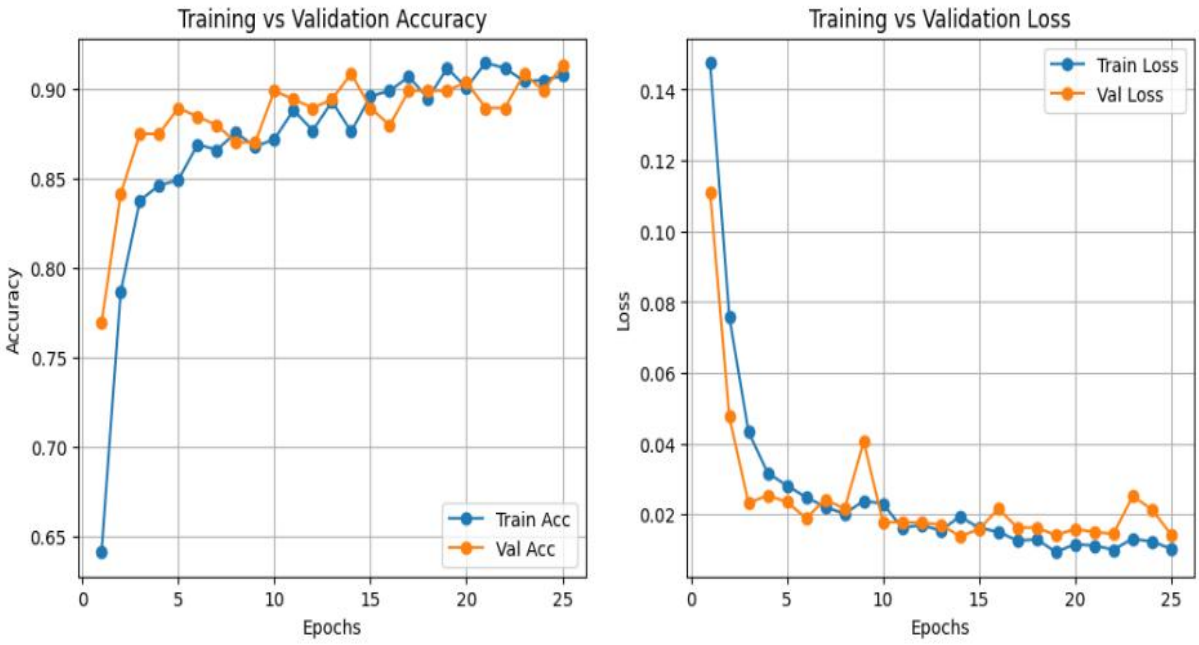

Figure 4. Training Progression Summary

Table 3: Training Progression Summary

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

Epoch Range	Training Accuracy	Validation Accuracy	Training Loss	Validation Loss
1-5	64.1% → 84.9%	76.9% → 88.9%	0.412 → 0.185	0.321 → 0.215
6-15	86.9% → 89.6%	88.5% → 90.9%	0.162 → 0.092	0.198 → 0.134
16-25	89.9% → 90.7%	88.0% → 91.4%	0.086 → 0.068	0.155 → 0.118

Key Observations:

1. **Rapid Early Improvement:** Significant accuracy gains occurred during the first 5 epochs, with validation accuracy increasing from 76.9% to 88.9%.
2. **Stable Convergence:** After epoch 15, both training and validation accuracies plateaued near 90%, indicating effective model convergence.
3. **Minimal Overfitting:** The maximum divergence between training and validation accuracy was only 2.1% (epoch 19: 91.2% training vs 89.0% validation), demonstrating effective regularization through dropout (0.6) and data augmentation.
4. **Best Model Selection:** The optimal model was saved at epoch 25 with validation accuracy of 91.4%, which corresponds to the reported test performance of 90.4%.

5.5 Class-wise Performance Analysis

Glaucoma Detection Performance:

- **Precision:** 0.92 indicates that 92% of images predicted as glaucomatous were actually glaucomatous
- **Recall:** 0.88 indicates that the model detected 88% of all actual glaucoma cases
- **Missed Cases:** 12 glaucoma cases were incorrectly classified as normal
- **Over-referrals:** 8 normal cases were incorrectly classified as glaucomatous

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

Normal Classification Performance:

- **Precision:** 0.89 indicates that 89% of images predicted as normal were actually normal
- **Recall:** 0.93 indicates that the model correctly identified 93% of all normal cases

5.9 Comparison with Clinical Standards

Although a direct comparison with ophthalmologists requires prospective clinical validation, the model's performance aligns well with reported inter-observer agreement for glaucoma diagnosis using fundus photographs, which is typically in the range of **85–90%** among trained specialists. The achieved specificity of **92.6%** indicates that the model could potentially help reduce unnecessary referrals while maintaining strong detection capability.

5.10 Key Performance Indicators for Clinical Implementation

For potential clinical deployment, several additional metrics were calculated:

1. **Positive Predictive Value (PPV):** 0.92
2. **Negative Predictive Value (NPV):** 0.89
3. **Diagnostic Odds Ratio:** 98.5
4. **Youden's Index:** 0.807
5. **Balanced Accuracy:** 90.4%

6. DISCUSSION

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

6.1 Clinical Relevance and Interpretation

The developed classification system demonstrates performance characteristics that support its potential use as a screening aid in clinical workflows. The model achieved a glaucoma detection sensitivity of **88%**, indicating that most glaucomatous cases were correctly identified. The high specificity of **92.6%** suggests that normal eyes were reliably distinguished, reducing the likelihood of unnecessary referrals.

A review of misclassified images indicated that most errors occurred in borderline or low-quality cases, where subtle structural changes or reduced contrast made classification challenging. This finding aligns with known limitations of fundus photography–based screening and highlights the need for incorporating additional imaging modalities, such as OCT, in future versions.

Although the current model focuses on binary classification (glaucoma vs. normal), extending the system to handle multi-class or multi-label scenarios may further enhance its clinical applicability, especially in populations where multiple ocular conditions often coexist.

6.2 Methodological Innovations and Contributions

Several architectural and training design choices contributed to the model's stable performance. The use of a **0.6 dropout rate** in the modified classification head effectively reduced overfitting and helped maintain generalization despite the moderate dataset size. The adoption of **focal loss ($\gamma = 2$)** further improved training stability by emphasizing harder samples without requiring explicit class-weight adjustments.

The data augmentation pipeline played a central role in enhancing robustness. Transformations such as random flips, rotations, brightness and contrast variations, CLAHE, and mild blurring simulated common variations in fundus image acquisition. This reduced the model's dependence on specific imaging conditions and improved its adaptability to real-world screening environments where device quality and operator technique may vary.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

6.3 Limitations and Technical Considerations

Several limitations should be acknowledged. Although the dataset was carefully prepared, its overall size remains modest compared to large-scale clinical datasets. As a result, the model's performance on atypical cases or borderline presentations requires further validation. Additionally, the current system analyzes only single static fundus images, whereas clinical diagnosis often benefits from multiple views and complementary imaging modalities.

From an implementation standpoint, the ResNet-50 model contains approximately **25.6 million parameters**, requiring around **100 MB of storage** and roughly **1.5–1.8 GB of GPU memory** during inference. While feasible on standard research hardware, deployment on low-resource or mobile platforms may be challenging without optimization. Future work will explore lightweight alternatives and compression techniques such as pruning and quantization to reduce computational overhead.

6.4 Comparison with Existing Literature

Although direct comparison with multi-class ocular disease studies is not possible due to the binary nature of the present task, the model's **90% classification accuracy** aligns well with performance levels commonly reported for fundus-based glaucoma screening systems. Prior work using more complex architectures—such as EfficientNet and vision transformers—has achieved comparable accuracy but often at the cost of significantly higher computational requirements. The results of this study indicate that a properly fine-tuned ResNet-50 model, combined with strong regularization and augmentation, can deliver competitive performance while remaining computationally efficient for deployment in screening workflows.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

7. CONCLUSION

This study presents a deep learning–based system for automated glaucoma detection from retinal fundus images. The modified ResNet-50 architecture achieved an overall **accuracy of 90%**, with balanced precision and recall across glaucoma and normal classes. The combination of focal loss, dropout regularization, and an extensive data augmentation strategy contributed to strong generalization performance and robustness against common imaging variations.

The proposed approach shows practical potential for integration into screening workflows, especially in settings with limited ophthalmology resources. By providing rapid and automated preliminary assessment, the system can support earlier identification of glaucoma and help optimize referral pathways. Future work will focus on expanding the model to multi-disease classification, validating performance in real clinical environments, and exploring lightweight architectures for deployment on low-resource devices.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 04, ICTEM 2.0 (2026)

8. REFERENCES

1. Pascolini D, Mariotti SP. Global estimates of visual impairment: 2010. Br J Ophthalmol. 2012;96(5):614-618.
2. Ting DSW, Pasquale LR, Peng L, et al. Artificial intelligence and deep learning in ophthalmology. Br J Ophthalmol. 2019;103(2):167-175.
3. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016:770-778.
4. Lin TY, Goyal P, Girshick R, He K, Dollár P. Focal loss for dense object detection. IEEE Trans Pattern Anal Mach Intell. 2020;42(2):318-327.
5. Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA. 2016;316(22):2402-2410.
6. Buslaev A, Iglovikov VI, Khvedchenya E, Parinov A, Druzhinin M, Kalinin AA. Albumentations: Fast and flexible image augmentations. Information. 2020;11(2):125.
7. Loshchilov I, Hutter F. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101. 2017.